

WHAT IS CLAIMED IS:

1. A method of distributing a message from a source to one or more hosts in a communications network that comprises one or more multicast capable routers, the method comprising:

5 in a first multicast capable router, generating a first table that identifies sub-branches connected to the first multicast capable router, and for each identified sub-branch, includes membership-related information from which it can be determined how many hosts are part of the sub-branch;

10 in the first multicast capable router, receiving the message to be distributed and a first value, *N1*, that specifies how many hosts the message is to be distributed to;

in the first multicast capable router, using the first table to identify *N1* hosts and their respective sub-branches; and

15 in the first multicast capable router, distributing the message to each of the identified sub-branches along with a second value, *N2*, that specifies how many hosts within the sub-branch the message is to be distributed to.

2. The method of claim 1, wherein:

at least one of the identified sub-branches includes a second multicast capable router; and

20 the method further comprises:

in each of the second multicast capable routers, generating a second table that identifies sub-branches connected to the second multicast capable router, and for each identified sub-branch, includes membership-related information from which it can be determined how many hosts are part of the sub-branch;

25 in the second multicast capable router, using the second table to identify *N2* hosts and their respective sub-branches; and

in the second multicast capable router, distributing the message to each of the identified sub-branches along with a third value, $N3$, that specifies how many hosts within the sub-branch the message is to be distributed to.

3. The method of claim 1, wherein:

the first table further includes, for each identified sub-branch, a first metric associated with each of the sub-branch hosts; and

the step of using the first table to identify NI hosts and their respective sub-branches includes, for each table entry, examining the first metric to determine whether a host satisfies criteria for being selected as one of the NI hosts.

4. The method of claim 3, wherein:

the table further includes at least one additional metric; and

the step of using the first table to identify NI hosts and their respective sub-branches further includes, for each table entry, examining the at least one additional metric to determine whether a host satisfies criteria for being selected as one of the NI hosts.

5. The method of claim 3, wherein each first metric indicates a distance between the sub-branch host and the second multicast capable router.

6. The method of claim 5, wherein the step of using the first table to identify NI hosts and their respective sub-branches comprises using the first table to identify NI nearest hosts and their respective sub-branches.

7. The method of claim 3, wherein each first metric indicates a delay between the sub-branch host and the second multicast capable router.

8. The method of claim 3, wherein each first metric indicates a cost associated with communication between the sub-branch host and the second multicast capable router.

9. The method of claim 3, wherein each first metric indicates a bandwidth associated with communication between the sub-branch host and the second multicast capable router.

10. The method of claim 3, wherein each first metric indicates a transmission delay between the sub-branch host and the second multicast capable router.

11. The method of claim 10, further comprising:
in the first multicast router, determining a transmission delay between the first multicast router and the sub-branch host.

12. The method of claim 3, wherein each first metric indicates a reliability of communication between the sub-branch host and the second multicast capable router.

13. The method of claim 1, wherein at least one of the identified sub-branches includes a router that does not have multicast capability.

14. The method of claim 1, wherein the step of generating the first table comprises:

sending a membership query message from the first multicast capable router to each sub-branch connected to the first multicast capable router;
receiving a membership information response from each sub-branch in response to the membership query message, wherein each membership information

response identifies, for the sub-branch, an identification of each multicast capable router in the sub-branch; and

generating the first table from the membership information response received from the sub-branches.

5 15. The method of claim 14, wherein each membership information response further identifies, for the sub-branch, a corresponding first metric for said each host in the sub-branch.

10 16. The method of claim 14, wherein each multicast capable router in each sub-branch responds to receipt of the membership query message by:
 determining a delay time; and
 sending a membership information response after the delay time has elapsed only if said each multicast capable router in said each sub-branch has not detected a membership information response being sent by a different multicast capable router in said each sub-branch.

15 17. The method of claim 1, wherein the first table is generated periodically.

 18. The method of claim 1, wherein:
 at least one of the identified sub-branches includes a second multicast capable router; and
 the method further comprises:
20 in each of the second multicast capable routers, unicasting the message to one or more hosts within the sub-branch.

 19. The method of claim 18, wherein from each of the second multicast capable routers, the message is unicast to N^2 hosts within the sub-branch.

20. The method of claim 1, wherein:

at least one of the identified sub-branches includes a second multicast capable router; and

the method further comprises:

5 in each of the second multicast capable routers, multicasting the message to all hosts within the sub-branch.

21. The method of claim 20, further comprising:

10 in each of the second multicast capable routers, first determining from the message whether it is permissible to distribute the message to more than $N2$ hosts prior to multicasting the message to all hosts within the sub-branch.

22. The method of claim 21, wherein an indication of whether it is permissible to distribute the message to more than $N2$ hosts is included in an optional header extension of the message.

15 23. A method of distributing a message from a source to one or more hosts in a communications network that comprises one or more multicast capable routers, the method comprising:

20 in a first multicast capable router, generating a first table that identifies sub-branches connected to the first multicast capable router, and for each identified sub-branch, includes membership-related information from which it can be determined how many hosts are part of the sub-branch;

in the first multicast capable router, receiving the message to be distributed and a first value, $N1$, that specifies a percentage of all multicast group members that the message is to be distributed to;

25 in the first multicast capable router, using the first table to identify a set of hosts that constitute $N1$ percent of all multicast group member hosts and their respective sub-branches; and

in the first multicast capable router, distributing the message to each of the identified sub-branches along with a second value, $N2$, that specifies how many hosts within the sub-branch the message is to be distributed to.

24. A method of determining how many hosts are members of a sub-branch that is connected to a router in a digital communications network, the method comprising:

 sending a membership query message from the router to the sub-branch;
 and

 receiving a membership information response from the sub-branch in response to the membership query message, wherein the membership information response identifies, for the sub-branch, an identification of each multicast capable router in the sub-branch and a value representing the number of hosts connected to said each multicast capable router.

25. The method of claim 24, wherein each multicast capable router in the sub-branch responds to receipt of the membership query message by:

 determining a delay time; and

 sending a membership information response after the delay time has elapsed only if said each multicast capable router in said each sub-branch has not detected a membership information response being sent by a different multicast capable router in said each sub-branch.

26. The method of claim 24, wherein the membership information further identifies, for the sub-branch, a corresponding first metric for each said host connected to said each multicast capable router in the sub-branch.

27. The method of claim 26, wherein the first metric indicates a distance between the host and the multicast capable router.

28. The method of claim 24, wherein each sub-branch recursively generates a membership information response by propagating a membership query message downstream to each of the multicast capable routers within the sub-branch, and updating a membership information response as it is returned upstream from one multicast capable router to the next.

29. The method of claim 24, further comprising generating the value representing the number of hosts connected to said each multicast capable router by:

multicasting membership information from each member in the sub-branch to a local router of the sub-branch and to all other hosts connected to local network.

30. The method of claim 24, further comprising generating the value representing the number of hosts connected to said each multicast capable router by:

unicasting membership information from each member in the sub-branch to a local router of the sub-branch; and

generating the membership information response in the local router, based on membership information received from said each member in the sub-branch.

31. An apparatus for distributing a message from a source to one or more hosts in a communications network that comprises one or more multicast capable routers, the apparatus comprising:

in a first multicast capable router, logic that generates a first table that identifies sub-branches connected to the first multicast capable router, and for each identified sub-branch, includes membership-related information from which it can be determined how many hosts are part of the sub-branch;

in the first multicast capable router, logic that receives the message to be distributed and a first value, *N1*, that specifies how many hosts the message is to be distributed to;

in the first multicast capable router, logic that uses the first table to identify *N1* hosts and their respective sub-branches; and

in the first multicast capable router, logic that distributes the message to each of the identified sub-branches along with a second value, *N2*, that specifies how many hosts within the sub-branch the message is to be distributed to.

32. The apparatus of claim 31, wherein:

at least one of the identified sub-branches includes a second multicast capable router; and

the apparatus further comprises:

in each of the second multicast capable routers, logic that generates a second table that identifies sub-branches connected to the second multicast capable router, and for each identified sub-branch, includes membership-related information from which it can be determined how many hosts are part of the sub-branch;

in the second multicast capable router, logic that uses the second table to identify *N2* hosts and their respective sub-branches; and

in the second multicast capable router, logic that distributes the message to each of the identified sub-branches along with a third value, *N3*, that specifies how many hosts within the sub-branch the message is to be distributed to.

33. The apparatus of claim 31, wherein:

the first table further includes, for each identified sub-branch, a first metric associated with each of the sub-branch hosts; and

the logic that uses the first table to identify *N1* hosts and their respective sub-branches includes logic that examines, for each table entry, the first metric to

determine whether a host satisfies criteria for being selected as one of the *NI* hosts.

34. The apparatus of claim 33, wherein:

the table further includes at least one additional metric; and

the logic that uses the first table to identify *NI* hosts and their respective sub-branches includes, logic that examines, for each table entry, the first metric and the at least one additional metric to determine whether a host satisfies criteria for being selected as one of the *NI* hosts.

35. The apparatus of claim 33, wherein each first metric indicates a distance between the sub-branch host and the second multicast capable router.

36. The apparatus of claim 35, wherein the logic that uses the first table to identify *NI* hosts and their respective sub-branches comprises logic that uses the first table to identify *NI* nearest hosts and their respective sub-branches.

37. The apparatus of claim 33, wherein each first metric indicates a delay between the sub-branch host and the second multicast capable router.

38. The apparatus of claim 33, wherein each first metric indicates a cost associated with communication between the sub-branch host and the second multicast capable router.

39. The apparatus of claim 33, wherein each first metric indicates a bandwidth associated with communication between the sub-branch host and the second multicast capable router.

40. The apparatus of claim 33, wherein each first metric indicates a transmission delay between the sub-branch host and the second multicast capable router.

41. The apparatus of claim 40, further comprising:

5 in the first multicast router, logic that determines a transmission delay between the first multicast router and the sub-branch host.

42. The apparatus of claim 33, wherein each first metric indicates a reliability of communication between the sub-branch host and the second multicast capable router.

10 43. The apparatus of claim 31, wherein at least one of the identified sub-branches includes a router that does not have multicast capability.

44. The apparatus of claim 31, wherein the logic that generates the first table comprises:

15 logic that sends a membership query message from the first multicast capable router to each sub-branch connected to the first multicast capable router;

 logic that receives a membership information response from each sub-branch in response to the membership query message, wherein each membership information response identifies, for the sub-branch, an identification of each multicast capable router in the sub-branch; and

20 logic that generates the first table from the membership information response received from the sub-branches.

45. The apparatus of claim 44, wherein each membership information response further identifies, for the sub-branch, a corresponding first metric for said each host in the sub-branch.

46. The apparatus of claim 44, wherein each multicast capable router in each sub-branch includes logic that responds to receipt of the membership query message, said logic comprising:

logic that determines a delay time; and

5 logic that sends a membership information response after the delay time has elapsed only if said each multicast capable router in said each sub-branch has not detected a membership information response being sent by a different multicast capable router in said each sub-branch.

47. The apparatus of claim 31, wherein the first table is generated periodically.

10 48. The apparatus of claim 31, wherein:

at least one of the identified sub-branches includes a second multicast capable router; and

the apparatus further comprises:

15 in each of the second multicast capable routers, logic that unicasts the message to one or more hosts within the sub-branch.

49. The apparatus of claim 48, wherein from each of the second multicast capable routers, the message is unicast to $N/2$ hosts within the sub-branch.

50. The apparatus of claim 31, wherein:

20 at least one of the identified sub-branches includes a second multicast capable router; and

the apparatus further comprises:

in each of the second multicast capable routers, logic that multicasts the message to all hosts within the sub-branch.

51. The apparatus of claim 50, further comprising:

logic that sends a membership query message from the router to the sub-branch; and

logic that receives a membership information response from the sub-branch in response to the membership query message, wherein the membership information response identifies, for the sub-branch, an identification of each multicast capable router in the sub-branch and a value representing the number of hosts connected to said each multicast capable router.

55. The apparatus of claim 54, wherein each multicast capable router in the sub-branch includes logic that responds to receipt of the membership query message, said logic comprising:

logic that determines a delay time; and

logic that sends a membership information response after the delay time has elapsed only if said each multicast capable router in said each sub-branch has not detected a membership information response being sent by a different multicast capable router in said each sub-branch.

56. The apparatus of claim 54, wherein the membership information further identifies, for the sub-branch, a corresponding first metric for each said host connected to said each multicast capable router in the sub-branch.

57. The apparatus of claim 56, wherein the first metric indicates a distance between the host and the multicast capable router.

58. The apparatus of claim 54, wherein each sub-branch includes logic that recursively generates a membership information response by propagating a membership query message downstream to each of the multicast capable routers within the sub-branch, and updates a membership information response as it is returned upstream from one multicast capable router to the next.

59. The apparatus of claim 54, further comprising logic that generates the value representing the number of hosts connected to said each multicast capable router by:

5 multicasting membership information from each member in the sub-branch to a local router of the sub-branch and to all other hosts connected to local network.

60. The apparatus of claim 54, further comprising generating the value representing the number of hosts connected to said each multicast capable router by:

10 unicasting membership information from each member in the sub-branch to a local router of the sub-branch; and

 generating the membership information response in the local router, based on membership information received from said each member in the sub-branch.